



Channel Charting for Streaming CSI Data

Sueda Taner, M. Guillaud, Olav Tirkkonen, Christoph Studer

► To cite this version:

Sueda Taner, M. Guillaud, Olav Tirkkonen, Christoph Studer. Channel Charting for Streaming CSI Data. Asilomar Conference on Signals, Systems, and Computers, Oct 2023, Pacific Grove, CA, United States. hal-04325927

HAL Id: hal-04325927

<https://inria.hal.science/hal-04325927>

Submitted on 6 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright

Channel Charting for Streaming CSI Data

Sueda Taner¹, Maxime Guillaud², Olav Tirkkonen³, and Christoph Studer¹

¹*Department of Information Technology and Electrical Engineering, ETH Zurich, Switzerland*

²*Inria & CITI Laboratory, France*

³*Department of Communications and Networking, Aalto University, Finland*

email: taners@iis.ee.ethz.ch, maxime.guillaud@inria.fr, olav.tirkkonen@aalto.fi, and studer@ethz.ch

Abstract—Channel charting (CC) applies dimensionality reduction to channel state information (CSI) data at the infrastructure basestation side with the goal of extracting pseudo-position information for each user. The self-supervised nature of CC enables predictive tasks that depend on user position without requiring any ground-truth position information. In this work, we focus on the practically relevant *streaming* CSI data scenario, in which CSI is constantly estimated. To deal with storage limitations, we develop a novel streaming CC architecture that maintains a small core CSI dataset from which the channel charts are learned. Curation of the core CSI dataset is achieved using a min-max-similarity criterion. Numerical validation with measured CSI data demonstrates that our method approaches the accuracy obtained from the complete CSI dataset while using only a fraction of CSI storage and avoiding catastrophic forgetting of old CSI data.

I. INTRODUCTION

Channel charting (CC), proposed in [1], applies dimensionality reduction to large sets of estimated channel state information (CSI) data acquired over long time periods. For a typical (fixed) wireless infrastructure access point (AP) or basestation (BS) and a fixed scattering environment (buildings etc.), CSI is predominantly influenced by the scattering environment. Thus, low-dimensional embeddings of such CSI datasets of mobile users produced by CC are expected to be topologically similar to certain features of the environment. This property has been verified experimentally, for instance in [2], where the two-dimensional embedding of the mobile users is topologically similar to the users' position on the ground. Here, topological similarity means that the obtained representation is accurate up to transformations such as rotating, stretching, and other continuous (but possibly nonlinear) transforms. Such transformations exist due to the self-supervised nature of CC, which also precludes its use for classical positioning. Nonetheless, CC associates a *pseudo position* to each CSI sample, which can be used for a broad range of location-based tasks—see [3] for an overview of CC and its applications.

Typical applications of CC involve repeatedly performing dimensionality reduction (DR) on large amounts of CSI samples that are continuously collected at high rates (e.g., tens to hundreds times per second) by APs or BSs with limited storage and computing power. While DR attempts to capture the geometry of the complete distribution, since the beginning

of observing CSI samples, storing *all* CSI data will inevitably exhaust the memory available at the AP/BS. Thus, in practice, only a tiny fraction of the collected CSI data can be stored long-term. In addition, since successive CSI samples cannot be assumed drawn independently from this distribution, due to the slow evolution of the geometric propagation parameters (e.g., moving users), classical approaches that deal with streaming data, such as time windowing, potentially result in catastrophic forgetting [4]. Therefore, wireless communication systems that generate a constant stream of new CSI samples require an online method that maintains a fixed dataset size, while still enabling the extraction of high-quality channel charts and avoiding catastrophic forgetting.

A. Contributions

Since the CC literature routinely ignores the fact that CSI data is acquired perpetually in a streaming fashion, we introduce a new approach for CC with streaming CSI data and limited memory. Our method consists of a simple dynamic core dataset selection strategy that updates the set of CSI samples used for learning the channel charts. This strategy has the benefit of conceptually separating the distillation algorithm from the DR process, which is performed solely based on the core CSI memory contents; see Fig. 1 for an overview of our approach. We introduce a core CSI memory curation strategy that dynamically minimizes the maximum similarity across the stored CSI samples. Results with the measured CSI dataset from [5] demonstrate that the proposed min-max-similarity curation approach avoids catastrophic forgetting and, even with a small core CSI memory, achieves comparable performance as an idealistic baseline that has access to *all* CSI samples.

B. Relevant Prior Art

In the field of machine learning, the problem of selecting a subset of a dataset while preserving its information content (measured as its potential to be used as a training set yielding a prescribed accuracy) is known as core-set selection [6]. A more general approach, known as dataset condensation (or distillation), consists of synthesizing a new composite, smaller dataset (not necessarily a subset of the full dataset) based on linear combinations of the original samples (see e.g. [7]). Dataset condensation appears to be a valid approach to deal with streaming data (see, e.g., the results in [8], [9] which leverage

The authors acknowledge the support of the European CHIST-ERA program through the CHASER (channel charting as a service) project.

dataset condensation through a gradient matching approach for the purpose of continual learning). However, existing dataset condensation approaches (see [10] for an overview) appear to be limited to *supervised* learning problems, and their extension to self-supervised DR is not straightforward. In contrast, our proposed approach is tailored to CC, which is an inherently self-supervised learning problem.

II. CHANNEL CHARTING OVERVIEW

We start by introducing the system model and summarizing traditional CC from non-streaming CSI.

A. System Model

We consider a single-input multiple-output (SIMO) wireless communication system, in which one or multiple single-antenna user equipments (UEs) transmit pilots to one or multiple APs/BSs, with a total number of B receive antennas. We assume that transmission utilizes orthogonal frequency-division multiplexing (OFDM) with W used subcarriers. For one UE at position $\mathbf{x}^{(n)} \in \mathbb{R}^3$, the CSI vector estimated by the BS at the w th subcarrier is $\mathbf{h}_w^{(n)} \in \mathbb{C}^B$, where n denotes the sample index. The CSI matrix associated with the UE at position $\mathbf{x}^{(n)}$ is obtained by concatenating the channel vectors from all subcarriers $[\mathbf{h}_1^{(n)}, \dots, \mathbf{h}_W^{(n)}] = \mathbf{H}^{(n)} \in \mathbb{C}^{B \times W}$.

B. Channel Charting Basics

Traditional CC operates in the following two phases [1].

1) *Learning the CC Function*: In the first phase, a large CSI dataset from various UE positions is collected at one or many infrastructure basestations (BSs). To capture the large-scale fading characteristics of the wireless channel [1] and to improve resilience against system and hardware impairments [2], [11], one first transforms each CSI matrix $\mathbf{H}^{(n)}$ into a CSI feature $\mathbf{f}^{(n)} \in \mathbb{R}^{D'}$. We represent this transform by a feature extraction function $f: \mathbb{C}^{B \times W} \rightarrow \mathbb{R}^{D'}$, so that $\mathbf{f}^{(n)} = f(\mathbf{H}^{(n)})$. The total dataset of acquired CSI features is $\{\mathbf{f}^{(n)}\}_{n=1}^N$, where N denotes the total number of acquired CSI samples. We note that if CSI is acquired at fast rates and over long periods of time, then storing all $D' \times N$ scalars may be infeasible.

From the CSI feature dataset, one then learns a CC function $g_\theta: \mathbb{R}^{D'} \rightarrow \mathbb{R}^D$, where $D \ll D'$, so that each CSI feature $\mathbf{f}^{(n)}$ is mapped to a D -dimensional pseudo-position $\hat{\mathbf{x}}^{(n)} = g_\theta(\mathbf{f}^{(n)})$. Here, the vector θ denotes the trainable parameters of the dimensionality reduction function. Note that the CC function is trained in a self-supervised manner, solely from the CSI feature dataset $\{\mathbf{f}^{(n)}\}_{n=1}^N$, which does *not* use any ground-truth UE position information—training typically aims at preserving pairwise relationships (e.g., matching dissimilarities) between feature space and latent space. The resulting, low-dimensional latent-space representation is the *channel chart*.

2) *Channel Charting*: In the second phase, the learned CC function g_θ is then used for channel charting, where one maps (typically in real time) new CSI features to pseudo-positions in the channel chart. The benefit of a channel chart is that it captures pseudo-positions of the transmitting UEs by preserving the local geometry: UEs that are close in space are also close in the channel chart.

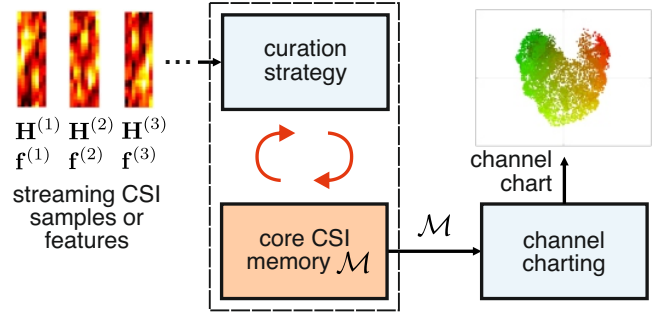


Fig. 1. Illustration of channel charting for streaming CSI data. A curation strategy maintains a subset of the obtained CSI samples or features in a small core CSI memory \mathcal{M} , which is used to perform channel charting.

III. CHANNEL CHARTING WITH STREAMING CSI

We now present our architecture for CC from streaming CSI and then discuss two core memory curation strategies.

A. Architecture Overview

In a realistic scenario, APs or BSs continuously estimate CSI matrices at high rates, which prevents them from storing all CSI data over long periods of time. We address this problem by learning a CC function from streaming CSI under the constraint of limited memory at the BSs. Concretely, we propose an architecture as depicted in Fig. 1. Here, either CSI matrices or features arrive perpetually and in a streaming fashion. A distillation algorithm, which is responsible for curating a core CSI dataset \mathcal{M} that stores no more than M CSI samples, assesses each newly-arrived CSI sample and either stores it in the core CSI memory (which may require one to replace a datapoint in the core CSI memory) or discards it.¹ The CC function g_θ is then learned from the core CSI dataset \mathcal{M} .

B. Core CSI Memory Curation Strategies

The core CSI memory is then curated by storing only a subset of the streaming CSI data (matrices or features) as follows. At the beginning, all of the estimated CSI samples are added to the core memory until its capacity M is reached. After capacity has been reached, we assess each new CSI sample. If we decide to add the new CSI sample to the core CSI memory, then we maintain fixed memory capacity by replacing one existing CSI sample. Otherwise, the memory remains unchanged and the new CSI sample is discarded. We now propose two curation strategies that determine whether to add a new CSI sample and which existing CSI sample to discard.

1) *Random Subset (RandoS)*: The first (and simplest) curation strategy that would come to one's mind is to store a random subset of the streaming CSI. Concretely, after core memory capacity has been reached, this curation strategy randomly decides whether to update the core memory by adding a new CSI sample $\mathbf{H}^{(n)}$ with probability $p_{\text{update}} \in [0, 1]$. If one decides to add the new CSI sample, then one chooses to randomly

¹One could also directly store CSI features in the core CSI memory \mathcal{M} . In this work, we store the CSI matrices as it enables the use of elaborate dissimilarity metrics for channel charting; see Sec. IV-B for the details.

Algorithm 1 Similarity-based Subset (SimS) Curation Strategy

```

1: initialize  $\mathcal{M} = \emptyset$  and  $n = 1$ 
2: for each newly arrived CSI sample  $\mathbf{H}^{(n)}$  do
3:   if  $n \leq M$  then
4:      $\mathcal{M} \leftarrow \mathcal{M} \cup \mathbf{H}^{(n)}$ 
5:   else
6:      $s \leftarrow \max_{\mathbf{M}_i \in \mathcal{M}} \text{sim}(\mathbf{H}^{(n)}, \mathbf{M}_i)$ 
7:      $\{k, \ell\} \leftarrow \arg \max_{1 \leq i \neq j \leq M} \text{sim}(\mathbf{M}_i, \mathbf{M}_j)$ 
8:     if  $s < \text{sim}(\mathbf{M}_k, \mathbf{M}_\ell)$  then
9:       with probability  $p$ ,  $r \leftarrow k$ ; otherwise,  $r \leftarrow \ell$ 
10:       $\mathcal{M} \leftarrow (\mathcal{M} \setminus \mathbf{M}_r) \cup \mathbf{H}^{(n)}$ 
11:    end if
12:  end if
13:   $n \leftarrow n + 1$ 
14: end for

```

replace the r th element $\mathbf{H}^{(r)} \in \mathcal{M}$ according to the probability mass function (PMF) $p_r \in [0, 1]$ with $\sum_r p_r = 1$. One can choose p_{update} and the PMF p_r , $r = 1, \dots, M$, according to some preference (e.g., updating the core memory more frequently or replacing more recent CSI samples).

2) *Similarity-Based Subset (SimS)*: The second curation strategy takes into account the CC application and attempts to store a subset of the streaming CSI that minimizes the maximum similarity in the core memory. The intuition behind this strategy is to avoid redundancy by maintaining a core CSI memory \mathcal{M} with maximally dissimilar CSI samples. The implementation details of this curation strategy are given in Alg. 1 and the method proceeds as follows. For each new arriving CSI sample $\mathbf{H}^{(n)}$, after memory capacity M is reached, we first calculate the similarity between the new CSI sample and all of the CSI samples in the core memory (line 6). By denoting the i th element in the core CSI memory \mathcal{M} by \mathbf{M}_i , we find the pair from \mathcal{M} that exhibits maximum similarity $\{\mathbf{M}_k, \mathbf{M}_\ell\}$; this set contains the two candidates of samples that should be replaced (line 7). If the maximum of the similarities between the new CSI sample and the CSI samples in the core memory is smaller than the similarity between \mathbf{M}_k and \mathbf{M}_ℓ (line 8), then, with probability p , we replace \mathbf{M}_k with the new sample $\mathbf{H}^{(n)}$; otherwise, we replace \mathbf{M}_ℓ (lines 9 and 10). In our implementation, we measure the pairwise CSI similarity using the absolute cosine similarity between CSI features:

$$\text{sim}(\mathbf{H}^{(i)}, \mathbf{H}^{(j)}) = \frac{|(\mathbf{f}^{(i)})^H \mathbf{f}^{(j)}|}{\|\mathbf{f}^{(i)}\| \|\mathbf{f}^{(j)}\|}. \quad (1)$$

Here, $\mathbf{f}^{(i)} = f(\mathbf{H}^{(i)})$, the superscript H is the conjugate transpose, and $\|\cdot\|$ is the Euclidean norm.

IV. EXPERIMENTAL SETUP

We now describe our experimental setup, which we use in Sec. V to assess the efficacy of the proposed methods.

A. Dataset Description

We use the measured CSI from the DICHASUS dataset [5], [12]. These measurements are for an indoor distributed SIMO

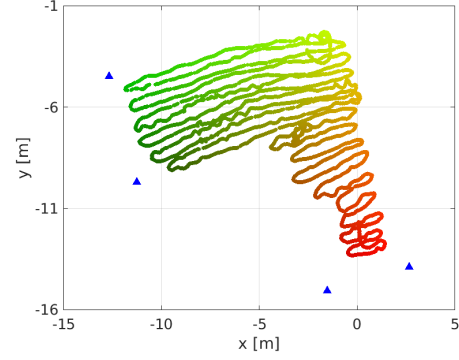


Fig. 2. Ground-truth UE positions (green-to-red gradient-colored area) of the streaming CSI and AP positions (blue triangles). Only a subset of the CSI samples stemming from these locations will be stored in the core memory \mathcal{M} .

communication system, in which a moving robot transmits pilots to four APs with eight antennas each. The APs are synchronized in frequency, time, and phase via over-the-air synchronization, and the transmitting robot uses an omnidirectional antenna. The system uses OFDM with a total number of $W = 1024$ subcarriers over a bandwidth of 50 MHz, centered at a carrier frequency of 1.272 GHz. The dataset in [12] comprises CSI from multiple trajectories of the robot over the same area.

To imitate a streaming CSI scenario, we use only the first trajectory from [12]. This trajectory yields a total number of $N = 17516$ sequentially-arriving CSI samples, where the robot follows a path from the top-left corner of the area to the bottom-right²; see Fig. 2 for the robot's trajectory.

B. Channel Charting Pipeline

The specifics of our channel charting pipeline are as follows. We adopt the CSI feature extraction function from [13]–[15], where we first apply an inverse discrete Fourier transform over the W subcarriers to transform the CSI into the delay domain. We denote the delay domain CSI vector for the n th UE position at delay tap τ with $\mathbf{h}_\tau^{(n)} \in \mathbb{C}^B$. We only keep the first C columns of the delay-domain CSI, $\{\mathbf{h}_\tau^{(n)}\}_{\tau=1}^C$, since the first few taps usually contain most of the received power. We vectorize this truncated-delay domain CSI matrix and take entrywise absolute values. Finally, we scale the resulting vector to unit Euclidean norm. The resulting CSI feature vectors $\mathbf{f}^{(n)} \in \mathbb{R}^{BC}$ will be the input of the CC function.

The literature describes a variety of neural-network-based channel charting functions building on autoencoders [1], [16], networks trained with a triplet loss [2], [15], [17], [18], or Siamese neural networks [14], [19], [20]. In what follows, we focus on Siamese neural networks that match the distances between the positions in the channel chart to pairwise dissimilarities computed from CSI. This property is expressed by the

²For easier interpretation of the generated results, we exclude the last 1 000 CSI samples where the robot moves back towards the start.

following loss:

$$\mathcal{L}(\theta) = \sum_{i=1}^{N-1} \sum_{j=i+1}^N (d_{i,j} - \|g_{\theta}(\mathbf{f}^{(i)}) - g_{\theta}(\mathbf{f}^{(j)})\|)^2. \quad (2)$$

Here, $d_{i,j}$ denotes the dissimilarity between $\mathbf{H}^{(i)}$ and $\mathbf{H}^{(j)}$ used for channel charting. The measure of dissimilarity we use is calculated in two steps: First, we compute the angle-delay profile (ADP)-based metric from [20], which is given by

$$\tilde{d}_{i,j} = \sum_{\tau=1}^C \left(1 - \frac{|(\mathbf{h}_{\tau}^{(i)})^H \mathbf{h}_{\tau}^{(j)}|^2}{\|\mathbf{h}_{\tau}^{(i)}\|^2 \|\mathbf{h}_{\tau}^{(j)}\|^2} \right). \quad (3)$$

Intuitively, this metric calculates the squared cosine similarity between two delay domain CSI vectors per-tap, and sums over the taps³. Second, we compute geodesic dissimilarities as in [19], [20]. To this end, we form a K -nearest neighbor graph of every sample according to its ADP-based dissimilarities with all other CSI samples. We then apply Dijkstra’s algorithm [21] on this graph to determine all shortest paths. Finally, the geodesic dissimilarity $d_{i,j}$ is given by the length of the shortest path between samples i and j .

The CC function g_{θ} is implemented using the neural network architecture from [2]. We train a six-layer fully-connected neural network with the following numbers of activations per layer: {256, 128, 64, 32, 16, 2}. All layers except the last one use ReLU activations; the last one uses linear activations. We use Glorot’s method [22] for weight initialization, and we utilize Adam [23] to train the channel charting function g_{θ} .

C. Core-Memory Size and Algorithm Parameters

We assume that the core memory capacity is $M = 1000$ and we store a subset of the streaming CSI using the two curation strategies proposed in Sec. III-B. We refer to the core CSI memories obtained with RandoS and SimS by \mathcal{M}_R and \mathcal{M}_S , respectively, to avoid ambiguity when we discuss one of them specifically. For RandoS, we set the probability $p_{\text{update}} = 0.5$ and $p_r = 1/M$ for all $\mathbf{H}^{(r)} \in \mathcal{M}_R$. For SimS, we set $p = 0.5$.

D. Performance Metrics

We assess the effectiveness of the proposed methods using four standard metrics from the CC literature; more details can be found in [24]. (i) Trustworthiness (TW) penalizes false neighborhood relationships in the channel chart, i.e., points that are neighbors in the channel chart but not in real world coordinates. (ii) Continuity (CT) quantifies preservation of real-world coordinate neighborhood relationships in the channel chart. (iii) Kruskal stress (KS) characterizes the mismatch between pairwise distances in real-world coordinates and those in the channel chart. (iv) Rajsiki distance (RD) measures the discrepancy between mutual information and joint entropy of the distribution of pairwise distances in the real-world coordinates and channel chart. All metrics range from 0 to 1. The optimal value for TW and CT is 1 (large is good); the optimal value for KS and RD is 0 (small is good).

³In case of a distributed scenario with multiple APs, one would use the CSI vectors per AP and sum over the APs.

TABLE I
CHANNEL CHARTING PERFORMANCE COMPARISON.

Method	Figure	Latent space quality metrics			
		TW \uparrow	CT \uparrow	KS \downarrow	RD \downarrow
RandoS	5 (b)	0.834	0.884	0.441	0.942
SimS	5 (c)	0.963	0.963	0.212	0.814
All	5 (d)	0.975	0.975	0.197	0.799

V. RESULTS

We are finally ready to demonstrate the efficacy of our approach to streaming CSI data.

A. Evolution of Core CSI Memory Contents

Figures 3 and 4 show the ground-truth positions corresponding to the CSI in the core memory after $n = 6000$, $n = 12000$, and $n = 17516$ arrived CSI samples. The color scheme represents the maximum cosine similarity from (1) of each sample with all others in the core memory, i.e., the color of UE position $\mathbf{x}^{(n)}$ is determined by $\max_{i \in \mathcal{M}} \text{sim}(\mathbf{H}^{(n)}, \mathbf{H}^{(i)})$.

In Fig. 3, we see that the core memory \mathcal{M}_R formed by RandoS experiences catastrophic forgetting—the memory stores mostly recent (and thus, similar) CSI data and old CSI samples are lost. In Fig. 4, we see that the core memory \mathcal{M}_S formed by SimS avoids catastrophic forgetting and contains CSI samples with dissimilar cosine similarities. We note that the maximum similarity values are generally lower for SimS in Fig. 4(c) than for the earlier time instants in (a) and (b), which is to be expected as the algorithm gradually reduces the maximum cosine similarity within the core memory over time. However, the maximum similarities of some samples seem to be larger in Fig. 4(b) than in (a); this shows that, although the SimS curation strategy aims at reducing the maximum similarity among all samples, the per-sample similarity may increase.

B. Channel Charting from Streaming CSI Data

We now show channel charts obtained from the proposed methods. We train one Siamese neural network from \mathcal{M}_R and one from \mathcal{M}_S after $N = 17516$ arrived CSI samples. As a baseline, we also train a Siamese network with *all* N arrived samples, which we refer as “All.” We then evaluate the channel chart quality performance using the second trajectory from [12]; this test set consists of 23478 CSI samples.

In Fig. 5, we show the ground-truth positions of the UE in the test set along with three channel charts. Fig. 5(b) confirms the expectation that RandoS does not perform well due to catastrophic forgetting. In fact, the training set \mathcal{M}_R includes almost no samples from the green part of the original trajectory. In stark contrast, Figs. 5(c) and (d) obtained from SimS and the idealistic “All” baseline that had the full CSI dataset for training, respectively, show high-quality channel charts in which the green and red areas are well separated.

To quantify the improvement of Sims over RandoS, and to compare Sims to the baseline “All,” Tbl. I lists the associated performance metrics outlined in Sec. IV-D. As expected,

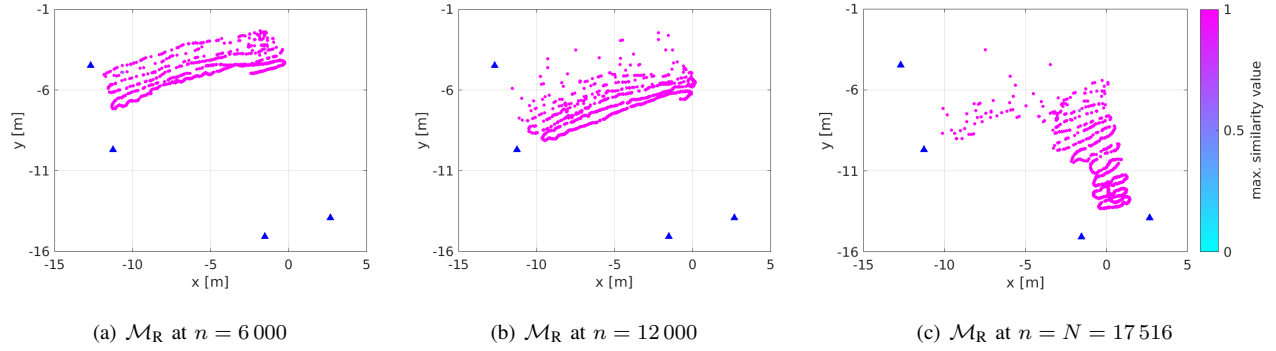


Fig. 3. Core memory evolution for RandoS. We show ground truth UE positions (cyan-to-pink gradient-colored area) corresponding to the CSI samples stored in the core memory \mathcal{M}_R after $n = 6000$, $n = 12000$, and $n = 17516$ arrived samples. The colors are determined by the maximum similarity of each sample with the other samples in the memory; triangles designate AP positions.

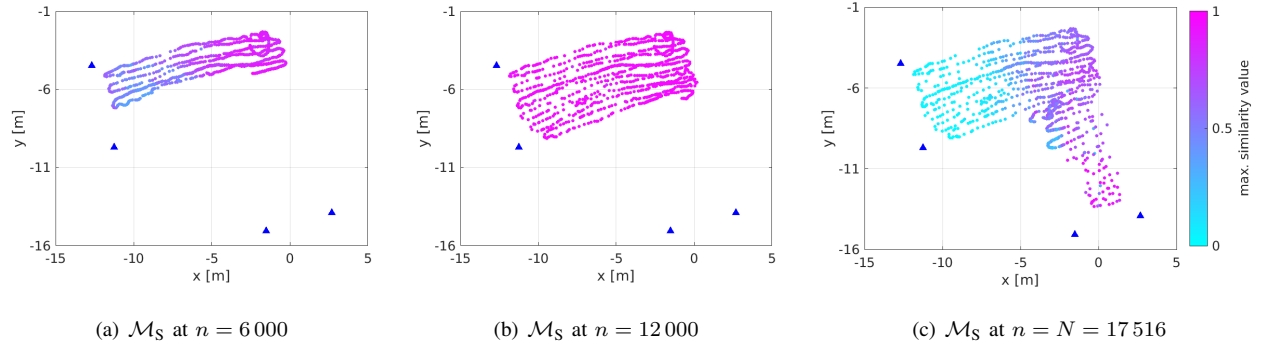


Fig. 4. Core memory evolution for SimS. We show ground truth UE positions (cyan-to-pink gradient-colored area) corresponding to the CSI samples stored in core memory \mathcal{M}_S after $n = 6000$, $n = 12000$, and $n = 17516$ arrived samples. The colors are determined by the maximum cosine similarity of each CSI sample with all the other samples in the core memory; triangles designate AP positions.

RandoS performs the worst in all considered metrics and the “All” baseline performs best. Quite surprisingly, the performance of SimS with a very small core memory of only $M = 1000$ CSI samples is at most 0.015 worse than the “All” baseline in all four metrics. Put simply: learning a CC function from a tiny core CSI memory curated with the SimS strategy achieves comparable quality as training from *all* 17516 CSI samples!

VI. CONCLUSIONS AND FUTURE WORK

We have studied a realistic scenario in which CSI is estimated in a streaming fashion and the basestation has a small and fixed-capacity core memory that stores CSI data from which one learns the channel charting function. We have proposed two simple curation strategies to maintain the core memory: The first one stores a random subset of the streaming CSI and is mostly used to demonstrate catastrophic forgetting. The second utilizes a criterion that attempts to minimize the maximum cosine similarity between CSI samples in the core memory. We have demonstrated with measured CSI data that our first method indeed suffers from catastrophic forgetting, while the second method results in a core memory that contains a representative subset of the streamed CSI data. The latter approach with a core dataset of only 1000 CSI samples results in comparable

channel chart quality as an idealistic approach that has access to all 17516 CSI samples.

There are many avenues for future work. Other similarity metrics and more sophisticated curation strategies should be explored. Furthermore, the core CSI memory could store data that is more suitable for channel charting (e.g., linear combinations of arrived CSI features). Finally, the consistency of channel charts over time should be studied.

REFERENCES

- [1] C. Studer, S. Medjkouh, E. Gönültaş, T. Goldstein, and O. Tirkkonen, “Channel charting: Locating users within the radio environment using channel state information,” *IEEE Access*, vol. 6, pp. 47 682–47 698, Aug. 2018.
- [2] P. Ferrand, A. Decurninge, L. G. Ordoñez, and M. Guillaud, “Triplet-based wireless channel charting: Architecture and experiments,” *IEEE J. Sel. Areas Commun.*, vol. 39, no. 8, pp. 2361–2373, Jun. 2021.
- [3] P. Ferrand, M. Guillaud, C. Studer, and O. Tirkkonen, “Wireless channel charting: Theory, practice, and applications,” *IEEE Commun. Mag.*, vol. 61, no. 6, pp. 124–130, Jun. 2023.
- [4] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska *et al.*, “Overcoming catastrophic forgetting in neural networks,” *Proc. National Academy of Sciences*, vol. 114, no. 13, pp. 3521–3526, Mar. 2017.
- [5] F. Euchner, M. Gauger, S. Dörner, and S. ten Brink, “A distributed massive MIMO channel sounder for “big CSI data”-driven machine learning,” in *Proc. Int. ITG Workshop on Smart Antennas (WSA)*, Nov. 2021.

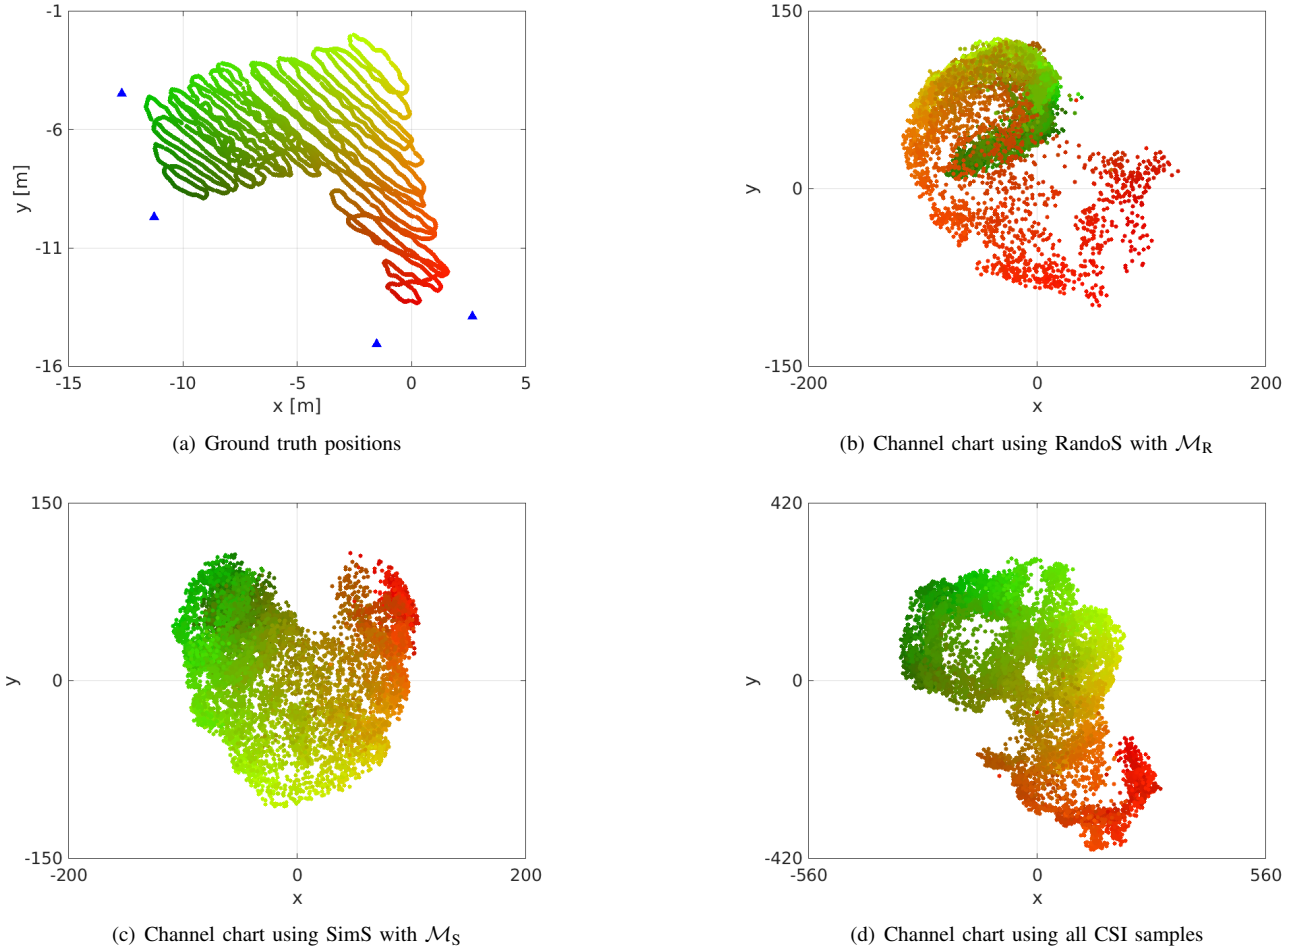


Fig. 5. Ground truth positions (a) and channel charts (b-d) for the test set. In (b) and (c), the channel charting-performing neural networks are trained from the core CSI memories \mathcal{M}_R and \mathcal{M}_S , respectively. In (d), all $N = 17\,516$ samples of streaming CSI are used for CC function training.

- [6] B. Mirzasoleiman, J. Bilmes, and J. Leskovec, "Coresets for data-efficient training of machine learning models," in *Proc. Int. Conf. on Mach. Learn.*, vol. 119, Jul. 2020, pp. 6950–6960.
- [7] B. Zhao, K. R. Mopuri, and H. Bilen, "Dataset condensation with gradient matching," *arXiv preprint arXiv:2006.05929*, Mar. 2021.
- [8] F. Wiewel and B. Yang, "Condensed composite memory continual learning," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, Jul. 2021.
- [9] M. Sangermano, A. Carta, A. Cossu, and D. Bacciu, "Sample condensation in online continual learning," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, Jul. 2022.
- [10] R. Yu, S. Liu, and X. Wang, "Dataset distillation: A comprehensive review," *IEEE Trans. Pattern Anal. Mach. Intell.*, Oct. 2023, early access.
- [11] E. Gönültaş, E. Lei, J. Langerman, H. Huang, and C. Studer, "CSI-based multi-antenna and multi-point indoor positioning using probability fusion," *IEEE Trans. Wireless Commun.*, vol. 21, no. 4, pp. 2162–2176, Apr. 2022.
- [12] F. Euchner and M. Gauger, "CSI Dataset dichasus-cf0x: Distributed Antenna Setup in Industrial Environment, Day 1," 2022. [Online]. Available: <https://doi.org/doi:10.18419/darus-2854>
- [13] J. Vieira, E. Leitinger, M. Sarajlic, X. Li, and F. Tufvesson, "Deep convolutional neural networks for massive MIMO fingerprint-based positioning," in *Proc. IEEE Intl. Symp. Personal, Indoor, Mobile Radio Commun.*, Oct. 2017, pp. 1–6.
- [14] E. Lei, O. Castañeda, O. Tirkkonen, T. Goldstein, and C. Studer, "Siamese neural networks for wireless positioning and channel charting," in *Proc. Allerton Conf. Commun., Contr., Comput.*, Sep. 2019, pp. 200–207.
- [15] S. Taner, V. Palhares, and C. Studer, "Channel charting in real-world coordinates," presented at the IEEE Global Telecommun. Conf. (GLOBE-COM), Dec. 2023.
- [16] P. Huang, O. Castañeda, E. Gönültaş, S. Medjkouh, O. Tirkkonen, T. Goldstein, and C. Studer, "Improving channel charting with representation-constrained autoencoders," in *Proc. IEEE Int. Workshop Signal Process. Advances Wireless Commun. (SPAWC)*, Aug. 2019, pp. 1–5.
- [17] T. Yassine, L. L. Magoarou, S. Paquelet, and M. Crussière, "Leveraging triplet loss and nonlinear dimensionality reduction for on-the-fly channel charting," in *Proc. IEEE Int. Workshop Signal Process. Advances Wireless Commun. (SPAWC)*, Jul. 2022.
- [18] B. Rappaport, E. Gönültaş, J. Hoydis, M. Arnold, P. K. Srinath, and C. Studer, "Improving channel charting using a split triplet loss and an inertial regularizer," in *Int. Symp. on Wireless Commun. Sys. (ISWCS)*, Sep. 2021.
- [19] M. Stahlke, G. Yammine, T. Feigl, B. M. Eskofier, and C. Mutschler, "Indoor localization with robust global channel charting: A time-distance-based approach," *IEEE Trans. on Mach. Learn. in Commun. and Netw.*, Mar. 2023, early access.
- [20] P. Stephan, F. Euchner, and S. ten Brink, "Angle-delay profile-based and timestamp-aided dissimilarity metrics for channel charting," *arXiv preprint arXiv:2308.09539*, Sep. 2023.
- [21] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische mathematik*, vol. 1, no. 1, pp. 269–271, 1959.
- [22] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proc. Thirteenth Int. Conf. Artificial Intelligence Stat.*, vol. 9, May 2010, pp. 249–256.
- [23] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, Jan. 2017.
- [24] H. Al-Tous, P. Kazemi, C. Studer, and O. Tirkkonen, "Channel charting with angle-delay-power-profile features and earth-mover distance," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, Oct. 2022, pp. 1195–1201.